

A very fast inference algorithm for finite-dimensional spin glasses: Belief Propagation on the dual lattice.

Alejandro Lage-Castellanos, Roberto Mulet
*Department of Theoretical Physics, Physics Faculty,
University of Havana, La Habana, CP 10400, Cuba.*

Federico Ricci-Tersenghi
*Dipartimento di Fisica, INFN – Sezione di Roma 1 and CNR – IPCF, UOS di Roma,
Università La Sapienza, P.le A. Moro 5, 00185 Roma, Italy*

Tommaso Rizzo
Dipartimento di Fisica, Università La Sapienza, P.le A. Moro 5, 00185 Roma, Italy
(Dated: January 12, 2013)

Starting from a Cluster Variational Method, and inspired by the correctness of the paramagnetic Ansatz (at high temperatures in general, and at any temperature in the 2D Edwards-Anderson model) we propose a novel message passing algorithm — the Dual algorithm — to estimate the marginal probabilities of spin glasses on finite dimensional lattices. We show that in a wide range of temperatures our algorithm compares very well with Monte Carlo simulations, with the Double Loop algorithm and with exact calculation of the ground state of 2D systems with bimodal and Gaussian interactions. Moreover it is usually 100 times faster than other provably convergent methods, as the Double Loop algorithm.

I. INTRODUCTION

Inference problems are common to almost any scientific discipline. Often an inference problem can be recast in that of computing some marginal probability on a small subset of variables, given a joint probability distribution over a large number N of variables, $\{s_1, \dots, s_N\} \equiv \sigma$. A typical example is the computation of the expectation value of a variable $\langle s_i \rangle \equiv \sum_{s_i} s_i p_i(s_i)$ where $p_i(s_i)$ is the single-variable marginal probability defined as

$$p_i(s_i) \equiv \sum_{\sigma \setminus s_i} P(\sigma) .$$

Clearly the computation of the sum in the r.h.s. is as difficult as the computation of the partition function $Z \equiv \sum_{\sigma} P(\sigma)$, which is the main subject of statistical mechanics. In the general (and the most interesting) case, this problem can not be solved exactly in a time growing sub-exponentially with the size N . We are deemed to use some kind of approximation in order to compute marginals in a time growing linearly in N . The approximation schemes used so far are mainly adopted from the field of statistical mechanics [1], where mean-field-like approximations are standard and well controlled tools for approximating the free-energy, $-\beta f = \ln Z$.

For non-disordered models, like the Ising ferromagnetic model, these approximations work quite well and provide good results [2]. However, much less is known for models with disorder, and for this reason we will focus on spin glasses in the present paper. In disordered models one usually deals with an ensemble of problems (e.g. in spin glasses each sample has its own couplings, randomly chosen from a given distribution), and the results obtained by the statistical mechanics tools refer to average quantities, i.e. those of the typical samples. In other words one is not concerned with the behavior of a specific sample, but rather one looks at the whole ensemble. On the contrary when doing inference one is interested on the properties of a single specific problem and thus the above approximation schemes (based on statistical mechanics mean-field approaches) need to be converted into an algorithm that can be run on such a specific sample. Computing marginals on a given sample clearly gives more information than computing averages over the ensemble.

To our knowledge, effective (i.e. linear time in N) algorithms for computing marginals can be essentially divided in two broad classes: stochastic local search algorithms, that roughly sample the configurational space according to $P(\sigma)$, and algorithms based on some kind of mean-field approximation. The former are exact on the long run, but the latter can be much more useful if an approximated answer is required in a short time. Unfortunately the latter also have some additional drawback due to the mean-field nature of the underlying approximation, e.g. the appearance of spurious phase transitions, that may prevent the proper convergence of the algorithm.

One more reason why this latter class of algorithms has a limited scope of application is that the convergence of the algorithm may strongly depend on the presence of short loops in the network of variables interactions. In this sense the successful application of one of these algorithms to models defined on regular lattices (which have many short loops) would be a major achievement.

In this paper we introduce a fast algorithm for computing marginals in 2D and 3D spin glass models [3] defined by the Hamiltonian (further details are given below)

$$\mathcal{H}(\sigma) = - \sum_{(i,j)} J_{ij} s_i s_j . \quad (1)$$

The first non trivial mean-field approximation for the above model corresponds to the Bethe-Peierls approximation scheme and the Belief Propagation (BP) algorithm [4]. Unfortunately when BP is run on a given spin glass sample defined on a D dimensional lattice, it seems to provide exactly the same output as if it were run on a random regular graph with fixed degree $2D$: that

is, for $T \gtrsim T_{\text{Bethe}}$ it converges to a solution with all null local marginals ($\langle s_i \rangle = 0$), while for $T \lesssim T_{\text{Bethe}}$ it does not converge to a fixed point.

The next step in the series of mean-field approximations (also known as Kikuchi approximations or ‘cluster variation method’) is to consider joint probability distributions of the four spins belonging to the same plaquette [2, 5]. Under this approximation an algorithm has been derived which is called Generalized Belief Propagation (GBP) [6]. To our knowledge, this algorithm has been applied to 2D spin glasses, but only in presence of an external magnetic field, which is known to improve the convergence properties of GBP. Our experience says that, running plain GBP on a generic sample of a 2D spin glass without external field, a fixed point is reached only for high enough temperatures, well above the interesting region. The lack of convergence of GBP (and other similar message-passing algorithms, MPA) is a well known problem, whose solution is in general far from being understood. For this reason, a new class of algorithms has been recently introduced [7], which provably converge to a fixed point: these algorithms use a double loop iterative procedure (to be compared to the single loop in GBP) and for this reason they are usually quite slow.

The algorithm we are going to introduce is as fast as BP (which is the fastest algorithm presently available), and converges in a wider range of temperatures than BP. Moreover the marginal probabilities provided by our algorithm are as accurate as those which can be obtained by Double Loop algorithms in a much larger running time.

Our algorithm works in the absence of external field, that is in the situation where MPA have more problems in converging to a fixed point. In this sense it is a very important extension to presently available inference algorithms.

The rest of the work is organized as follows. First, in section II we derive the GBP equations for the 2D Edward Anderson model. Then, in section II A we rewrite these equations in terms of fields, a notation that has a nicer physical interpretation and that we are going to use in the rest of the work. Section III presents the novel algorithm, inspired by the paramagnetic Ansatz to the GBP equations. In section IV we show the results of running this algorithm on the 2D Edward-Anderson model. There we compare its performance with Monte Carlo simulations, with the Double Loop algorithm and with exact calculation of the ground state of systems with bimodal and Gaussian interactions. Then, in section V we generalized our message passing equation to general dimensions and present some results for the 3D Edward-Anderson model. Finally, some conclusions are drawn in section VI.

II. GENERALIZED BELIEF PROPAGATION ON THE 2D EA MODEL

Here we present the GBP equations for the Edwards-Anderson (EA) model on a 2D square lattice, and we refer the reader to [2, 6] for a more general introduction. In our case (as well as in many other cases) GBP is equivalent to Kikuchi’s approximation, known as the Cluster Variational Method (CVM) [5]. We will try a presentation as physical as possible.

Consider the 2D EA model consisting of a set $\sigma = \{s_1, \dots, s_N\}$ of N Ising spins $s_i = \pm 1$ located at the nodes of a 2D square lattice, interacting with a Hamiltonian

$$\mathcal{H}(\sigma) = - \sum_{\langle i,j \rangle} J_{ij} s_i s_j, \quad (2)$$

where the sum runs over all couples of neighboring spins (first neighbors on the lattice). The J_{ij} are the coupling constants between spins and are supposed to be fixed for any given instance of the model. If the interactions are not random variables, *i.e.* $J_{ij} = J$, then the 2D ferromagnet is recovered. We will focus on the two most common disorder distributions: bimodal interactions, $P(J) = \frac{1}{2}\delta(J-1) + \frac{1}{2}\delta(J+1)$, and Gaussian interactions $P(J) = \exp(-J^2/2)/\sqrt{2\pi}$.

The statistical mechanics of the EA model, at a given temperature $T = 1/\beta$, is given by the Gibbs-Boltzmann distribution

$$P(\sigma) = \frac{e^{-\beta\mathcal{H}(\sigma)}}{Z} .$$

The direct computation of the partition function

$$Z = \sum_{\sigma} e^{-\beta\mathcal{H}(\sigma)}$$

or any marginal distribution $p(s_i, s_j) = \sum_{\sigma \setminus s_i, s_j} P(\sigma)$ is a time consuming task, unattainable in practice, since it involves the addition of 2^N terms, and therefore an approximation is required¹.

The idea of the Region Graph Approximation to Free Energy [6] is to replace the real distribution $P(\sigma)$ by a reduced set of its marginals. The hierarchy of approximations is given by the size of such marginals, starting with the set of all single spins marginals $p_i(s_i)$ [mean-field], then following to all neighboring sites marginals $p_{ij}(s_i, s_j)$ [Bethe], then to all square plaquettes marginals $p_{ijkl}(s_i, s_j, s_k, s_l)$, and so on. Since the only way of knowing such marginals exactly is the unattainable computation of Z , the method pretends to approximate them by a set of beliefs $b_i(s_i)$, $b_{ij}(s_i, s_j)$, etc. obtained from the minimization of a region based free energy. In the Region Graph Approximation to Free Energy, a set of regions, *i.e.* sets of variables and their interactions, is defined, and a free energy is written in terms of the beliefs at each region. The Cluster Variational Method does a similar job, but instead of starting from an arbitrary choice of regions, it starts by defining the set of largest regions, and smaller regions are defined recursively by the intersections of bigger regions. In this sense, CVM is a specific choice of region graph approximation to the free energy.

For the 2D EA model, we will consider the expansion of the free energy in terms of the marginals at three levels of regions: single sites (or spins), links, and plaquettes. By plaquettes we mean the square basic cell of the 2D lattice. This choice of regions corresponds to the CVM having the square plaquettes as biggest regions. The free energy of the system is therefore written as

$$F = \sum_R c_R F_R ,$$

where R runs over all regions considered, and the free energy in a particular region depends on the marginals at that level $b_R(\sigma_R)$:

$$\beta F_R = \sum_{\sigma_R} b_R(\sigma_R) \beta E_R(\sigma_R) + b_R(\sigma_R) \log b_R(\sigma_R) .$$

The symbol σ_R refers to the set of spins in region R , while E_R is the energy contribution in that region. The counting numbers c_R (also Moebius coefficients) are needed to ensure that bigger regions do not over count the contribution in free energy of smaller regions, and follow the prescription

$$c_R = 1 - \sum_{R' \supset R} c_{R'} , \quad (3)$$

where R' is any region containing completely region R , as e.g. a plaquette containing a link or a link containing a site. In the case of the 2D lattice, the biggest regions are the square plaquettes,

¹ The 2D case is special: indeed, thanks to the small genus topology, the partition function Z can be computed efficiently. However we are interested in developing an algorithm for the general case, and we will not make use of this peculiarity.

and therefore $c_{\text{plaq}} = 1$, while the links regions have $c_{\text{link}} = 1 - 2c_{\text{plaq}} = -1$ (as each of them is contained in 2 plaquettes regions), and finally the spins regions have $c_{\text{site}} = 1 - 4c_{\text{plaq}} - 4c_{\text{link}} = 1$ (as each spin belongs to 4 links and 4 plaquettes). So the actual approximation for the EA model on a 2D square lattice is

$$\begin{aligned}
\beta F = & \sum_{\mathcal{P}} \sum_{\sigma_{\mathcal{P}}} b_{\mathcal{P}}(\sigma_{\mathcal{P}}) \log \frac{b_{\mathcal{P}}(\sigma_{\mathcal{P}})}{\exp[-\beta E_{\mathcal{P}}(\sigma_{\mathcal{P}})]} && \text{Plaquettes} \\
& - \sum_L \sum_{\sigma_L} b_L(\sigma_L) \log \frac{b_L(\sigma_L)}{\exp[-\beta E_L(\sigma_L)]} && \text{Links} \\
& + \sum_i \sum_{s_i} b_i(s_i) \log \frac{b_i(s_i)}{\exp[-\beta E_i(s_i)]} && \text{Sites}
\end{aligned} \tag{4}$$

where the sums run over all plaquettes, links and sites respectively. Please note that we are using the following notation for region indices: lower-case for sites, upper-case for links and upper-case calligraphic for plaquettes. The energy term $E_i(s_i)$ in the sites contribution is only relevant when an external field acts over spins, and can be taken as zero in our case, since no external field is considered. Notice that whenever the interactions are included in more than one region (in our case are included in Link and Plaquette regions), the counting numbers guarantee that the exact thermodynamical energy $U = \sum_{\sigma} P(\sigma) \mathcal{H}(\sigma)$ is obtained when the beliefs are the exact marginals of the Boltzmann distribution. On the other hand, the entropy contribution is intrinsically approximated, since the cutoff in the regions sizes imposes a certain kind of factorization of $P(\sigma)$ in terms of its marginals (see [2] for an explanation of the region graph approximation in terms of cumulants expansions of the entropy).

The next step in the method, is to compute the beliefs from the minimization condition of the free energy. However, an unrestricted minimization will generally produce inconsistent solutions, since the beliefs (marginals) are not independent, as they are related by the marginalization conditions

$$\begin{aligned}
b_i(s_i) &= \sum_{\sigma_{L \setminus i}} b_L(\sigma_L) = \sum_{s_j} b_L(s_i, s_j) , \\
b_L(\sigma_L) &= b_L(s_i, s_j) = \sum_{\sigma_{\mathcal{P} \setminus L}} b_{\mathcal{P}}(\sigma_{\mathcal{P}}) = \sum_{s_k, s_l} b_{\mathcal{P}}(s_i, s_j, s_k, s_l) ,
\end{aligned} \tag{5}$$

where $\sigma_L = \{s_i, s_j\}$ and $\sigma_{\mathcal{P}} = \{s_i, s_j, s_k, s_l\}$. In order to minimize under the constraints in Eq. (5) and under the normalization condition for each belief, a set of Lagrange multipliers should be added to the free energy in Eq. (4). There are different ways of choosing the Lagrange multipliers [6], and each of them will produce a different set of self consistency equations. We choose the so called Parent to Child scheme (see section XX in [6]), in which constraints in Eq. (5) are imposed by two sets of Lagrange multipliers: $\mu_{L \rightarrow i}(s_i)$ relating the belief at link L to that at site i , and $\nu_{\mathcal{P} \rightarrow L}(\sigma_L)$ relating the one at plaquette \mathcal{P} to the one at link L .

With constraints (5) enforced by Lagrange multipliers, the free energy stationary conditions for

the beliefs are the following:

$$\begin{aligned}
 b_i(s_i) &= \frac{1}{Z_i} \exp \left(-\beta E_i(s_i) - \sum_{L \supset i}^4 \mu_{L \rightarrow i}(s_i) \right), \\
 b_L(\sigma_L) &= \frac{1}{Z_L} \exp \left(-\beta E_L(\sigma_L) - \sum_{\mathcal{P} \supset L}^2 \nu_{\mathcal{P} \rightarrow L}(\sigma_L) - \sum_{i \in L}^2 \sum_{\substack{L' \supset i \\ L' \neq L}}^3 \mu_{L' \rightarrow i}(s_i) \right), \\
 b_{\mathcal{P}}(\sigma_{\mathcal{P}}) &= \frac{1}{Z_{\mathcal{P}}} \exp \left(-\beta E_{\mathcal{P}}(\sigma_{\mathcal{P}}) - \sum_{L \subset \mathcal{P}}^4 \sum_{\substack{\mathcal{P}' \supset L \\ \mathcal{P}' \neq \mathcal{P}}}^1 \nu_{\mathcal{P}' \rightarrow L}(\sigma_L) - \sum_{i \in \mathcal{P}}^4 \sum_{\substack{L \supset i \\ L \not\subset \mathcal{P}}}^2 \mu_{L \rightarrow i}(s_i) \right),
 \end{aligned} \tag{6}$$

where the notation $L \supset i$ refers to all links containing site i and $\mathcal{P} \supset L$ to all plaquettes containing link L . The upper indices in the sums are written just to help understanding how many terms are in each sum for the 2D case. The precise meaning of the indices in each summation can be understood from the graphical representation in Figure 1. Lagrange multipliers are shown as arrows going from parent regions to children regions: simple arrows correspond to $\mu_{L \rightarrow i}$ and triple arrows to $\nu_{\mathcal{P} \rightarrow L}$. Let us consider, for instance, the belief in a link region $b_L(\sigma_L)$, depicted in the central picture of Fig. 1: the sum of the two Lagrange multipliers $\nu_{\mathcal{P} \rightarrow L}(\sigma_L)$ corresponds to the triple arrows from plaquettes on the left and right of the link L , while the double sum over the three $\mu_{L \rightarrow i}(s_i)$ and the three $\mu_{L \rightarrow j}(s_j)$ correspond to the three arrows acting over the two spins.




FIG. 1. Schematic representation of belief equations (6). Lagrange multipliers are depicted as arrows, going from parent regions to children regions.

In Eq.(6), the Z_R are normalization constants, and the terms $E_{\mathcal{P}}(\sigma_{\mathcal{P}}) = E_{\mathcal{P}}(s_i, s_j, s_k, s_l) = -(J_{ij}s_i s_j + J_{jk}s_j s_k + J_{kl}s_k s_l + J_{li}s_l s_i)$ and $E_L(\sigma_L) = E_L(s_i, s_j) = -J_{ij}s_i s_j$ are the corresponding energies in plaquettes and links, and are represented in Fig. 1 by bold lines (interactions) between circles (spins). In our case $E_i(s_i)$ is zero since no external field is acting on the spins.

The Lagrange multipliers are fixed by the constraints they were supposed to enforce, Eq. (5),

and they must satisfy the following set of self-consistency equations:

$$\begin{aligned} \exp\left[-\mu_{L \rightarrow i}(s_i)\right] &= \sum_{s_j} \exp\left[-\beta E_{L \setminus i}(s_i, s_j) - \sum_{\mathcal{P} \supset L}^2 \nu_{\mathcal{P} \rightarrow L}(s_i, s_j) - \sum_{\substack{L' \supset j \\ L' \neq L}}^3 \mu_{L' \rightarrow j}(s_j)\right] \\ \exp\left[-\nu_{\mathcal{P} \rightarrow L}(s_i, s_j) - \mu_{D \rightarrow i}(s_i) - \mu_{U \rightarrow j}(s_j)\right] &= \\ \sum_{s_k, s_l} \exp\left[-\beta E_{\mathcal{P} \setminus L}(s_i, s_j, s_k, s_l) - \sum_{\substack{L' \in \mathcal{P} \\ L' \neq L}}^3 \sum_{\substack{\mathcal{P}' \supset L' \\ \mathcal{P}' \neq \mathcal{P}}}^1 \nu_{\mathcal{P}' \rightarrow L'}(\sigma_{L'}) - \sum_{\substack{L' \supset k \\ L' \not\subset \mathcal{P}}}^2 \mu_{L' \rightarrow k}(s_k) - \sum_{\substack{L' \supset l \\ L' \not\subset \mathcal{P}}}^2 \mu_{L' \rightarrow l}(s_l)\right] \end{aligned} \quad (7)$$

Again, to help understanding these equations, we provide in Fig. 2 their graphical representation. Note that there is one of these equations for every pair of Link-Site and every pair of Plaquette-Link in the graph. With $E_{P \setminus L}$ we refer to interactions in plaquette P that are not in link L .




FIG. 2. Message passing equations (7), shown schematically. Messages are depicted as arrows, going from parent regions to children regions. On any link J_{ij} , represented as bold lines between spins (circles), a Boltzmann factor $e^{\beta J_{ij} s_i s_j}$ exists. Dark circles represent spins to be traced over. Messages from plaquettes to links $\nu_{\mathcal{P} \rightarrow L}(s_i, s_j)$ are represented by a triple arrow, because they can be written in terms of three parameters U, u_i and u_j , defining the correlation $\langle s_i s_j \rangle$ and magnetizations $\langle s_i \rangle$ and $\langle s_j \rangle$, respectively.

For each link L in the 2D lattice, there are two link-to-site multipliers, $\mu_{L \rightarrow i}(s_i)$ and $\mu_{L \rightarrow j}(s_j)$. For each plaquette there are four plaquette-to-link multipliers $\nu_{\mathcal{P} \rightarrow L}(s_i, s_j)$, corresponding to the four links contained inside the plaquette. Let N be the number of spins in the lattice, there are $2N$ links and N plaquettes. So the originally intractable problem of computing marginals, has been replaced by the problem of solving a set of $4N + 4N$ coupled equations for the Lagrange multipliers as those in Eq. (7). Once these equations are solved, the approximation for the marginals is obtained from Eq. (6) for the beliefs, and all thermodynamic quantities are derived from them as in Eq. (4).

Minimizing a region graph approximation to free energy, as that in Eq. (4) with constraints Eq. (5), or equivalently solving the set of self-consistent equations in Eq. (7), is still a non trivial task. Let us consider two ways of doing it. The first method is the “direct” minimization of the constrained free energy, using a Double Loop algorithm [7]. This method is quite solid, since it guarantees convergence to an extremal point of the constrained free energy, but it may be very slow to converge. The second method, which is generally faster but is not guaranteed to converge, is the family of the so called Message-Passing algorithm (MPA), in which the Lagrange multipliers are

interpreted as *messages* $\nu_{\mathcal{P} \rightarrow L}(\sigma_L)$ going from plaquettes to links, and messages $\mu_{L \rightarrow i}(s_i)$ from links to sites. Self consistency equations (7) can be viewed as the update rules for the messages in the left hand side, in terms of those in the right hand side. A random order updating of the messages in the graph by Eq. (7) (message passing) can reach a fixed point solution, and therefore, to an extremal point of the constrained free energy [6]. Next, we show explicitly how the message-passing equations looks like in terms of fields.

A. From multipliers to fields

A particularly useful way of representing the multipliers (messages), with a nice physical interpretation, is the one used in [8], which we adopt here. In full generality [8, 9], these multipliers can be written in terms of effective fields

$$\mu_{L \rightarrow i}(s_i) = \beta u_{L \rightarrow i} s_i \quad (8)$$

$$\nu_{\mathcal{P} \rightarrow L}(s_i, s_j) = \beta (U_{\mathcal{P} \rightarrow L} s_i s_j + u_{\mathcal{P} \rightarrow i} s_i + u_{\mathcal{P} \rightarrow j} s_j) \quad (9)$$

In particular, the field u corresponds to the cavity field in the Bethe approximation [6]. Using Lagrange multipliers, messages or fields, is essentially equivalent. We will often refer to fields as u -messages to emphasize their role in a message-passing algorithm, and we will refer to self-consistency equations (7) as the message-passing equations.

This parametrization of the multipliers has proved useful to other endeavors, like the extension of the replica theory to general region graph approximations [9]. Here, all the relevant information in the Lagrange multipliers is translated to “effective fields” u and (U, u_a, u_b) . Notice that in this representation every single field u corresponds to an arrow in the schematic messages-passing equations in Figure 2. In particular, the messages going from plaquettes to links are characterized by three fields (U, u_a, u_b) , and the field U acts as an effective interaction term, that adds directly to the energy terms in the Boltzmann factor. For instance, the first message-passing Eq. (7) is

$$\exp [\beta u_{L \rightarrow i} s_i] = \sum_{s_j} \exp \left[\beta \left((u_{\mathcal{P} \rightarrow i} + u_{\mathcal{L} \rightarrow i}) s_i + (U_{\mathcal{P} \rightarrow L} + U_{\mathcal{L} \rightarrow L} + J_{ij}) s_i s_j + (u_{\mathcal{P} \rightarrow j} + u_{\mathcal{L} \rightarrow j} + u_{A \rightarrow j} + u_{B \rightarrow j} + u_{U \rightarrow j}) s_j \right) \right] \quad (10)$$

where the indices refer to the notation used in Fig. 2 and J_{ij} is the interaction coupling constant between spins s_i and s_j . This equation naturally defines the updating rule for the message $u_{L \rightarrow i}$:

$$u_{L \rightarrow i} = \hat{u}(u_{\mathcal{P} \rightarrow i} + u_{\mathcal{L} \rightarrow i}, U_{\mathcal{P} \rightarrow L} + U_{\mathcal{L} \rightarrow L} + J_{ij}, u_{\mathcal{P} \rightarrow j} + u_{\mathcal{L} \rightarrow j} + u_{A \rightarrow j} + u_{B \rightarrow j} + u_{U \rightarrow j}), \quad (11)$$

where

$$\hat{u}(u, U, h) \equiv u + \frac{1}{2\beta} \log \frac{\cosh \beta(U + h)}{\cosh \beta(U - h)}$$

Note that the usual cavity equation for fields in the Bethe approximation [10] is recovered if all contributions from plaquettes \mathcal{P} and \mathcal{L} are set to zero.

Working in a similar way for the second equation in (7) we end up with the updating rule for the message $(U_{\mathcal{P} \rightarrow L}, u_{\mathcal{P} \rightarrow i}, u_{\mathcal{P} \rightarrow j})$ sent from any given plaquette region \mathcal{P} to one of its children

links L (see right picture in Fig. 2)

$$\begin{aligned}
U_{\mathcal{P} \rightarrow L} &= \frac{1}{4\beta} \log \frac{K(1,1)K(-1,-1)}{K(1,-1)K(-1,1)} \\
u_{\mathcal{P} \rightarrow i} &= u_{\mathcal{D} \rightarrow i} - u_{D \rightarrow i} + \frac{1}{4\beta} \log \frac{K(1,1)K(1,-1)}{K(-1,1)K(-1,-1)} \\
u_{\mathcal{P} \rightarrow j} &= u_{\mathcal{U} \rightarrow j} - u_{U \rightarrow j} + \frac{1}{4\beta} \log \frac{K(1,1)K(-1,1)}{K(1,-1)K(-1,-1)}
\end{aligned} \tag{12}$$

where

$$\begin{aligned}
K(s_i, s_j) &= \sum_{s_k, s_l} \exp \left[\beta \left((U_{\mathcal{U} \rightarrow U} + J_{jk})s_j s_k + (U_{\mathcal{R} \rightarrow R} + J_{kl})s_k s_l + (U_{\mathcal{D} \rightarrow D} + J_{li})s_l s_i + \right. \right. \\
&\quad \left. \left. (u_{U \rightarrow k} + u_{C \rightarrow k} + u_{E \rightarrow k} + u_{R \rightarrow k})s_k + (u_{R \rightarrow l} + u_{F \rightarrow l} + u_{G \rightarrow l} + u_{D \rightarrow l})s_l \right) \right]
\end{aligned}$$

Equations (11) and (12) are equivalent to equations in (7), once multipliers (messages) are parametrized in terms of fields. For instance, note that the μ multipliers in the left hand side of second equation in (7) appear now subtracted in the right hand side of Eq. (12).

The field notation is more comprehensible and has a clear physical meaning. Each plaquette \mathcal{P} is telling its children links L that they should add an effective interaction term $U_{\mathcal{P} \rightarrow L}$ to the real interaction J_{ij} , due to the fact that spins s_i and s_j are also interacting through the other three links in the plaquette \mathcal{P} . Fields u act like magnetic fields upon spins, and the complete $\nu_{\mathcal{P} \rightarrow L}(s_i, s_j)$ -message is characterized by the triplet $(U_{\mathcal{P} \rightarrow L}, u_{\mathcal{P} \rightarrow i}, u_{\mathcal{P} \rightarrow j})$, and will be referred from now on as Uuu -message. Furthermore, it is clear that some fields enter directly in the message-passing equations like $u_{\mathcal{P} \rightarrow i}$ and $u_{\mathcal{L} \rightarrow i}$ in Eq. (11) and $u_{\mathcal{D} \rightarrow i}$ and $u_{\mathcal{U} \rightarrow j}$ in Eq. (12). Also note that since our model has no external field, the fields u break the symmetry of the original Hamiltonian whenever they are non zero. For instance, in the ferromagnetic model, when all $J_{ij} = J$, these fields are zero at high temperature and become non zero at Kikuchi's critical temperature $T = 1.42$ [5], implying a spontaneous magnetization in the ferromagnet.

III. THE DUAL APPROXIMATION FOR THE PARAMAGNETIC PHASE

Unfortunately, the iterative message-passing algorithm for solving the GBP equations (11) and (12) often does not converge on finite dimensional lattices. While this is expected if long range correlations are present, it is rather disappointing that it happens also in the paramagnetic phase, where one would like to find easily the solution to the model. Here we are going to focus only on the paramagnetic phase, and propose an improved solving algorithm based on physical assumptions.

In the paramagnetic phase of any spin model defined by the Hamiltonian in Eq. (2), that is with no external field, variables have no polarization or magnetization: this in turn implies that in the solution all u -fields must be zero, and only U -fields should be fixed self-consistently to non zero values.

This paramagnetic solution has some interesting properties. First, it is always a solution of the GBP equations, since Eq. (11) and Eq. (12) are self-consistent with all $u = 0$. This means that starting from unbiased messages (all $u = 0$) the iterative GBP algorithm keeps this property. Second, the paramagnetic Ansatz is correct, from the GBP perspective, at least at high enough temperatures, meaning that even if we start with biased messages ($u \neq 0$), the iterative algorithm converges to all $u = 0$ at high temperatures. And last, but not least, the well studied physical behavior for the 2D EA model with zero-mean random interactions J_{ij} , is expected to remain

always paramagnetic, *i.e.* to have no transition to a spin glass phase at any finite temperature [11]. Therefore, the Ansatz $u = 0$ is both physically plausible and algorithmically desirable.




FIG. 3. Message passing of correlation messages in the dual approximation. In the right hand side the trace is taken over the black spins.

Under the paramagnetic Ansatz, which we shall also call Dual approximation for a reason to be explained soon, the message-passing equation (11) is irrelevant, as it is always satisfied given that $\hat{u}(0, U, 0) = 0$, while Eq. (12) now turns into (see Fig. 3)

$$U_{\mathcal{P} \rightarrow L} = \hat{U}(U_{\mathcal{U} \rightarrow U}, U_{\mathcal{R} \rightarrow R}, U_{\mathcal{D} \rightarrow D}) = \frac{1}{\beta} \operatorname{arctanh} \left[\tanh \beta(U_{\mathcal{U} \rightarrow U} + J_{jk}) \tanh \beta(U_{\mathcal{R} \rightarrow R} + J_{kl}) \tanh \beta(U_{\mathcal{D} \rightarrow D} + J_{li}) \right]. \quad (13)$$

The only relevant messages now are those associated to the multipliers $\nu_{\mathcal{P} \rightarrow L}(s_i, s_j) = \beta U_{\mathcal{P} \rightarrow L} s_i s_j$, and they will be referred to as U -messages. Eq. (13) can be interpreted as a correlation message-passing equation, giving the new interaction field \hat{U} that a certain link shall experience as a consequence of the correlations transmitted around the plaquette. The belief Eq. (6) also simplifies. Obviously $b(s_i) = 0.5$ for every spin in the graph, and the link and plaquette beliefs are

$$b_L(s_i, s_j) = \frac{1}{Z_L} e^{\beta(U_{\mathcal{L} \rightarrow L} + U_{\mathcal{P} \rightarrow L} + J_{ij}) s_i s_j}, \quad (14)$$

$$b_{\mathcal{P}}(s_i, s_j, s_k, s_l) = \frac{1}{Z_{\mathcal{P}}} e^{\beta(U_{\mathcal{L} \rightarrow L} + J_{ij}) s_i s_j + \beta(U_{\mathcal{U} \rightarrow U} + J_{jk}) s_j s_k + \beta(U_{\mathcal{R} \rightarrow R} + J_{kl}) s_k s_l + \beta(U_{\mathcal{D} \rightarrow D} + J_{li}) s_l s_i}.$$

The Dual algorithm we are proposing to study the paramagnetic phase of the EA model, is a standard message passing algorithm for the U -messages, which works as follows.

- 1: Start with all U -messages null
- 2: **repeat**
- 3: Choose randomly one plaquette \mathcal{P} and one of its children links L
- 4: Update the field $U_{\mathcal{P} \rightarrow L}$ according to Eq. (13) as in Fig. 3
- 5: **until** The last change for any U -message is less than ϵ (we use typically $\epsilon = 10^{-10}$)
- 6: **return** The beliefs $b_L(s_i, s_j)$ defined in Eq. (14) for every pair of neighboring spins

Some damping factor $\gamma \in [0, 1)$ can be added in the update step $U_{\mathcal{P} \rightarrow L} = \gamma U_{\mathcal{P} \rightarrow L} + (1 - \gamma) \hat{U}$ in order to help convergence.

A. Mapping to the dual model

It is worth noticing that Eq. (13) is nothing but the BP equation for the corresponding dual model (hence the name of the algorithm).

The dual model has a binary variable $x_{ij} \equiv s_i s_j$ on every link of the original model, and the original coupling constants play now the role of an external polarizing (eventually random) field

$$\mathcal{H}_{\text{dual}}(\vec{x}) = - \sum_{\langle i,j \rangle} J_{ij} x_{ij} .$$

This Hamiltonian looks like the sum of independent variables, but this is not the case. The dual variables $x_{ij} = \pm 1$ must satisfy a constraint for each cycle (or closed path) in the original graph, enforcing that their product along the cycle must be equal to 1. On a regular lattice any closed path can be expressed in terms of elementary cycles of 4 links (the plaquettes) and so it is enough to enforce the constraint on every plaquette: $x_{ij} x_{jk} x_{kl} x_{li} = 1$. The Gibbs-Boltzmann probability distribution for the dual model is then given by

$$P(\vec{x}) = \frac{1}{Z} e^{-\beta \mathcal{H}_{\text{dual}}(\vec{x})} \prod_{\langle i,j,k,l \rangle} \delta_{x_{ij} x_{jk} x_{kl} x_{li}, 1} , \quad (15)$$

where the product runs over all elementary plaquette.

The model described by the probability measure in Eq. (15) can be viewed as a constraint satisfaction problem with a non uniform prior (given by $e^{-\beta \mathcal{H}_{\text{dual}}(\vec{x})}$). It is straightforward to derive the BP equations for such a problem. Indeed by defining the marginal for the variable x_{ij} on link L in the presence of the solely neighboring plaquette \mathcal{P} as $(1 + x_{ij} \tanh \beta U_{\mathcal{P} \rightarrow L})/2 \propto \exp(\beta U_{\mathcal{P} \rightarrow L} x_{ij})$, the BP equations read

$$\begin{aligned} & \frac{1}{2} (1 + x_{ij} \tanh \beta U_{\mathcal{P} \rightarrow L}) \propto \\ & \sum_{x_{jk}, x_{kl}, x_{li}} e^{\beta U_{\mathcal{U} \rightarrow \mathcal{U}} x_{jk}} e^{\beta J_{jk} x_{jk}} e^{\beta U_{\mathcal{R} \rightarrow \mathcal{R}} x_{kl}} e^{\beta J_{kl} x_{kl}} e^{\beta U_{\mathcal{D} \rightarrow \mathcal{D}} x_{li}} e^{\beta J_{li} x_{li}} \delta_{x_{jk} x_{kl} x_{li}, x_{ij}} \propto \\ & \sum_{\substack{x_{jk}, x_{kl}, x_{li}: \\ x_{jk} x_{kl} x_{li} = x_{ij}}} (1 + x_{jk} \tanh \beta (U_{\mathcal{U} \rightarrow \mathcal{U}} + J_{jk})) (1 + x_{kl} \tanh \beta (U_{\mathcal{R} \rightarrow \mathcal{R}} + J_{kl})) (1 + x_{li} \tanh \beta (U_{\mathcal{D} \rightarrow \mathcal{D}} + J_{li})) = \\ & 1 + x_{ij} \tanh \beta (U_{\mathcal{U} \rightarrow \mathcal{U}} + J_{jk}) \tanh \beta (U_{\mathcal{R} \rightarrow \mathcal{R}} + J_{kl}) \tanh \beta (U_{\mathcal{D} \rightarrow \mathcal{D}} + J_{li}) . \quad (16) \end{aligned}$$

In the second summation the terms containing one or two x variables sums to zero, while the other two terms are those written in the last expression. Equating the first and the last expressions, this equation is manifestly equal to Eq. (13).

B. Average Case Solution

GBP in general, and the Dual approximation in particular, are methods for the study of the thermodynamical properties of a given problem. However, in the limit of large systems ($N \rightarrow \infty$, thermodynamical limit), we expect a typical behavior to arise. This is the so called *self-averaging* property of disordered systems. By typical we mean that almost every realization of the interactions J_{ij} will result in a system whose thermodynamical properties (free energy, energy, entropy) are very close to the average value.

Normally, in disordered systems, we cope with the $N \rightarrow \infty$ limit and with the average over the random J_{ij} by the replica method. The application of the replica trick to regions graph

approximations is a challenging task [9]. However, we can still grasp the average case behavior with a cavity average case solution of the dual message-passing equations, at the price of neglecting the local structure of the graph (beyond plaquettes).

The idea is to represent the set of U -messages flowing in any given graph, by a population of messages $Q(U)$. Then the message-passing Eq. (13) is used to obtain such population in a self-consistent way. More precisely, in every iteration three messages U_1, U_2, U_3 are randomly drawn from the population $Q(U)$ and a new message $U_0 = \hat{U}(U_1, U_2, U_3)$ is computed by Eq. (13) using three couplings randomly selected from $P(J)$. The obtained message U_0 is put back into the population, and the iteration is repeated many times, until the population stabilizes.

Once we have the self consistent population of messages, we can compute the average energy

$$E_{\text{Ave}} = \langle -J_{ij} \tanh \beta (J_{ij} + U_1 + U_2) \rangle_{Q(U_1), Q(U_2), P(J_{ij})} \quad (17)$$

by a random sampling of the population and of the interactions. The average case solution is supposed to be very good whenever the network of interactions has no or few short loops. This is not the case in any finite dimensional lattice, since there the short loops (plaquettes) are abundant. Nonetheless, the average case solution gives a reasonably good approximation to the single instance results in 2D and 3D, as shown in the next Section.

IV. RESULTS ON 2D EA MODEL

Message-passing algorithms work fine in the high temperature regime ($T > T_c$) of models defined on random topologies: this is the reason why these methods have been successfully applied in random constraint satisfaction problems, like random-SAT or random-Coloring [12, 13]. However, when used on regular finite-dimensional lattices, they can experience difficulties even in the paramagnetic phase, because the presence of short loops spoils message-passing convergence.

It is well known that on a random graph of fixed degree (connectivity) $c = 4$ the cavity approximation gives a paramagnetic result above $T_{\text{Bethe}} \simeq 1.52$ (i.e. $\beta_{\text{Bethe}} \simeq 0.66$) with all cavity fields $u_i = 0$. Below the Bethe critical temperature, this solution becomes unstable to perturbations, and we expect many solutions to appear with non trivial messages $u_i \neq 0$. The presence of many solutions in the messages passing equations is connected to the existence of many thermodynamical states in the Gibbs-Boltzmann measure, or, equivalently, to the presence of replica symmetry breaking. The appearance of such a spin glass phase is also responsible for the lack of convergence of message-passing equations, since the intrinsic locality of the message-passing equations fails to coordinate distant regions of the graph (which are now long-range correlated). As a consequence, the application of BP to the 2D EA model (that also has fixed degree $c = 4$) still finds the paramagnetic phase at high temperatures, but below T_{Bethe} , the Bethe instability takes the message-passing iteration away from the $u = 0$ solution and does not allow the messages to convergence to a fixed point (i.e. the algorithm wanders forever). In Figure 4 we show the convergence probability for the BP message-passing equations in the 2D EA model.

On the other hand, a straightforward GBP Parent-to-Child implementation does not fully overcome this problem. At high temperatures, the Parent-to-Child equations converge to a paramagnetic solution with all $u = 0$ and non trivial $U \neq 0$, which turns out to be the same solution found by our Dual algorithm. When going down in temperature, the convergence properties of the algorithm worsen, and are sensitive to tricks like damping and bounding in the fields. A thorough discussion of these properties is left for a future work, but let us summarize that typically the algorithm stop converging at low temperatures, somewhere below T_{Bethe} , as shown in Figure 4.

So, in general, BP and GBP equations are not simple to use in finite dimensional systems at low enough temperatures: this warning was already reported in Refs. [2, 7, 14]. Indeed a




FIG. 4. Convergence probability of BP (Bethe approximation) and GBP on a 2D square lattice, as a function of inverse temperature. Data points are averages over 100 systems with random bimodal interactions. System sizes are $N = L^2$ with $L = 32, 128$ and a damping factor $\gamma = 0.5$ has been used in the iteration of the message-passing equations. The Bethe spin glass transition is expected to occur at $\beta_{\text{Bethe}} \simeq 0.66$ ($T_{\text{Bethe}} \simeq 1.52$) for a random graph with the same connectivity as the 2D square lattice. Notably, that temperature also marks the convergence threshold for BP equations in the 2D square lattice. GBP, on the contrary, reaches lower temperatures, but eventually stop converging.

different method for extremizing the constrained free energy named Double Loop algorithm [7, 15] was developed to overcome such difficulties. As mentioned earlier, Double Loop guarantees convergence of the beliefs, on any topology, with or without short loops. Given the convergence problems in GBP, researchers typically resort to Double Loop algorithms to extremize region graph approximations to the free energy, below the Bethe critical temperature.

In order to make a fair comparison with our Dual algorithm, we have used an optimized code for GBP and Double Loop algorithms: the open source LibDai library written in C++ [16].

The first interesting result of our work is that our Dual algorithm converges at all temperatures, just as Double Loop does. The reason why it converges is that there are no u -messages, so the Bethe instability will not affect our message-passing iteration.

The second relevant result of our Dual algorithm, is the fact that it finds the same solution found by the Double Loop algorithm at all temperatures. In other words, the direct extremization of the region graph approximation to free energy Eq. (4) via a Double Loop algorithm finds a paramagnetic solution characterized by the beliefs $b_i(s_i) = 0.5$ and $b_L(s_i, s_j) = \frac{1}{z} e^{-\beta \tilde{J}_{ij} s_i s_j}$; and the effective interactions \tilde{J}_{ij} found by the Double Loop algorithm are exactly equal to those found with our Dual algorithm, $\tilde{J}_{ij} = J_{ij} + U_{\mathcal{P} \rightarrow L} + U_{\mathcal{L} \rightarrow L}$. This means that beliefs and correlations found by the two algorithms are identical: $\langle s_i s_j \rangle_{\text{Double Loop}} = \langle s_i s_j \rangle_{\text{Dual}}$.

The third result is that the running times of our Dual algorithm are nearly four orders of magnitude smaller than those required by the Double Loop implementation in libdai, at least in a wide range of temperatures (see figure 5). More precisely, the convergence time of the Dual algorithm growth exponentially with $\beta = 1/T$, but still, in the relevant range of temperatures

where the region graph approximation is a good approximation (not too low temperatures), the running time is always roughly a factor 10^4 smaller than Double Loop.




FIG. 5. Running times of the Double Loop algorithm [7, 16] (libdai) and the Dual algorithm averaged over 10 realizations of a 2D 8×8 EA model with Gaussian interactions. Generally the Double Loop algorithm requires a time 4 orders of magnitude larger than that used by the Dual algorithm. Three different precision goals were used for the Dual algorithm 10^{-5} , 10^{-10} , 10^{-15} , while the precision of the Double Loop algorithm is 10^{-9} . The inset shows the behavior of the running times for both algorithm versus the system size $L = \sqrt{N}$.

A. Dual approximation *vs* Monte Carlo simulations

The fact that our Dual algorithm provides the same results (and much faster) than the Double Loop algorithm is a very good news. Essentially is telling us that we are not losing anything by restricting the space of possible messages, as far as the region graph approximation is concerned. However, the ultimate comparison for the approximation has to be done with the exact marginals and correlations. In figure 6 we show a comparison between the exact correlations $\langle s_i s_j \rangle_{\text{PT}}$ of neighboring spins obtained with a Parallel Tempering (PT) Monte Carlo simulation, and the Dual approximation estimate for the same two-spins correlations. The coincidence between $\langle s_i s_j \rangle_{\text{PT}}$ and $\langle s_i s_j \rangle_{\text{Dual}}$ is essentially perfect at high temperatures, and it becomes weaker as the temperature is decreased. The reason for the discrepancies is obviously the fact that we are using an approximation in which collective behaviors of spins is accounted exactly only until the plaquette level; more distant correlations are approximated and these correlations become more important at low temperatures.

Given such a good correspondence between the correlations under the Dual approximation and the true correlations, we expect a very good estimate for the energy too. In Figure 7 we show with points the energy under the Dual approximation and with full lines the Monte Carlo exact energy: the data are indeed very close. The dashed lines show the average case energy for the Dual




FIG. 6. Comparison between the correlations $\langle s_i s_j \rangle_{\text{Dual}}$ obtained by the Dual algorithm and the nearly exact correlations obtained by a Parallel Tempering simulation. We used a 64×64 EA model with Gaussian interactions. At each temperature the data correlation coefficient ρ is reported.

approximation, Eq. (17). In spite of the fact that the average case does not take into account the local structure of the lattice, the average case energy is quite close to the single instance one.

B. Ground State Configuration in 2D

The good agreement between the correlations found by the Dual algorithm, and those found in a Monte Carlo simulation, for the 2D EA model, compels us to push this correspondence down to $T = 0$. More precisely, using the correlations obtained by our Dual algorithm at low temperatures, we try to compute a ground state configuration by the following procedure. The idea is to freeze iteratively the relative position $s_i s_j$ of those interacting spins that are more strongly correlated, which is done by setting $J_{ij} \rightarrow \pm\infty$, and re-running the Dual algorithm until convergence every time one pair of spins is frozen. Note that freezing the relative position of spins is equivalent to freeze the dual variable $x_{ij} = s_i s_j$. The freezing procedure is very simple, but for the fact one has to check that frozen links must be consistent with a spin configuration. More precisely, frozen x_{ij} variables must satisfy the requirement that on any closed loops the product is one,

$$\prod_{ij \text{ closed loop}} x_{ij} = 1. \quad (18)$$

For very short loops the satisfaction of this condition is automatically induced by the Dual algorithm: for example if three links on a plaquette freeze, the fourth link is immediately frozen to a value satisfy condition in Eq. (18). However, for longer loops (as the one shown in Fig. 8), the propagation of these constraints by the Dual algorithm is not perfect, since the information




FIG. 7. Energy as a function of the inverse temperature β for a 64×64 2D EA model, with both types of interactions, Gaussian and bimodal. Full lines represent the exact thermodynamical energy as obtained by a Monte Carlo simulation, points are the energies obtained under the Dual approximation, and dashed lines are the average case energies.

degrades with distance beyond the plaquette level. Then we need to enforce the constraints of Eq. (18) by a proper algorithm. At each stage of the freezing process, we define the clusters of frozen links as follows: if two frozen links share a spin, then they belong to the same cluster. In Figure 8 a cluster of frozen links is represented by bold lines. Notice that, once a spin is fixed in a cluster, all other spins are fixed as well by the frozen correlations. On the other hand, different clusters of spins can have arbitrary relative orientations.




FIG. 8. Even if the link marked by the arrow is not the most polarized link according to the marginals provided by the Dual algorithm, the spins it connects are fully correlated by the fact that they belong to a cluster of frozen links (bold lines). Therefore, the marked link must be immediately fixed accordingly.

Consider now the situation depicted in Figure 8 and focus on the value of the correlation between the two spins connected by the link marked by an arrow. From the fact these two spins belong to the same cluster of frozen links (shown as bold link in Figure 8) we know they are perfectly

correlated, however by running the Dual algorithm we could get a weak value for this correlation and then proceed by freezing a different link. A set of sub-optimal choices of this kind may finally produce a configuration of frozen links where the constraints in Eq. (18) are not all satisfied. In order to avoid these constraint violations we force any link whose spins are already part of the same cluster to be polarized accordingly. The freezing algorithm, therefore, works as follows.

```

1: repeat
2:     run the Dual algorithm until convergence (at a low enough temperature)
3:     find the link  $L$  with largest finite  $\tilde{J}_L = J_L + U_{\mathcal{P} \rightarrow L} + U_{\mathcal{L} \rightarrow L}$ 
4:     freeze that link by setting  $J_L \leftarrow \text{Sign}(\tilde{J}_L) \infty$ 
5:     if link  $L$  is connected to clusters  $\mathbf{C}$  and  $\mathbf{C}'$  of frozen links then
6:         merge clusters  $\mathbf{C}$ ,  $\mathbf{C}'$  and link  $L$  in a unique cluster
7:     else
8:         if link  $L$  is connected to a single cluster  $\mathbf{C}$  of frozen links then
9:             add link  $L$  to cluster  $\mathbf{C}$ 
10:        else
11:            create a new cluster with link  $L$ 
12:        end if
13:    end if
14:    for all non-frozen links  $L'$  at the boundaries of a cluster of frozen links do
15:        if link  $L'$  shares both spins with the same cluster then
16:            freeze link  $L'$  accordingly {To avoid violations of constraint}
17:        end if
18:    end for
19: until all links are frozen
20: return the spins configuration obtained by setting one spin and fixing the rest according to
    frozen links

```

The results obtained with this freezing procedure are quite good. In figure 9 we compare the resulting ground state energies with the exact solutions obtained using a web service running an exact solving algorithm [17]. We used an ensemble of 100 EA models on the 2D square lattice with Gaussian interactions (so the ground state is not degenerate) and with bimodal interactions. Most points are along the bisecting line, meaning that the ground state found by either methods are the same. The relative error for the ground state energy is 0.0013 for Gaussian systems, and 0.00078 for bimodal systems. Looking at how many links are frustrated in one of the solutions and unfrustrated in the other, we found that the Dual+Freezing algorithm returns 94.3% of the correct link correlation signs with respect to the true ground state solution for the Gaussian models. For the bimodal system, given the degeneracy of the ground state configuration, it is more probable to find the actual ground state energy but for the same reason the link overlap between the exact Ground State and the one found with Dual+Freezing is significantly lower (86.0%).

In general we believe these results on the ground states to be very encouraging, considering that the Dual algorithm is very fast and not restricted to the 2D case (at variance to fast exact algorithms for computing ground states). They provide evidence that the marginals obtained by this Dual algorithm are reliable even at very low temperatures.

V. GENERALIZATION TO OTHER DIMENSIONS

Let us now consider the region graph based approximation to the free energy for a generic D -dimensional (hyper-)cubic lattice, using the same hierarchy of regions: square plaquettes, links and spins. After computing the counting numbers for a general D -dimensional lattice, see Eq. (3),




FIG. 9. Correlation of the Dual+Freezing ground state energy with the exact ground state energy in $N = 16 \times 16$ systems. The top left points correspond to 100 bimodal systems $J_{ij} = \pm 1$, while the right bottom points correspond to 100 systems with Gaussian interactions. For bimodal interactions, the degeneracy of the ground state improves the probability of actually finding the correct ground state energy (93%), and conversely reduces the expected link correlation overlap with the exact ground state solution (86%). For Gaussian interactions, the ground state is not degenerated, and only in $\sim 7\%$ of the samples the Dual+Freezing method finds the actual ground state; however the average link overlap is very high (94%). The line $f(x) = x$ is shown to guide the eye. Kindly note that two set of axes are being used.

the free energy approximation becomes

$$\begin{aligned}
 \beta F = & \sum_{\mathcal{P}} \sum_{\sigma_{\mathcal{P}}} b_{\mathcal{P}}(\sigma_{\mathcal{P}}) \log \frac{b_{\mathcal{P}}(\sigma_{\mathcal{P}})}{\exp(-\beta E_{\mathcal{P}}(\sigma_{\mathcal{P}}))} && \text{Plaquettes} \\
 & -(2D - 3) \sum_L \sum_{\sigma_L} b_L(\sigma_L) \log \frac{b_L(\sigma_L)}{\exp(-\beta E_L(\sigma_L))} && \text{Links} \\
 & +(2D^2 - 4D + 1) \sum_i \sum_{s_i} b_i(s_i) \log \frac{b_i(s_i)}{\exp(-\beta E_i(s_i))} && \text{Spins}
 \end{aligned} \tag{19}$$

Plaquettes are still the biggest regions considered at so have counting number 1, but now each link is contained in $2(D - 1)$ plaquettes, and each spin is in $2D$ links and $2D(D - 1)$ plaquettes. The message passing equations for the Dual algorithm in D dimensions are then

$$U_{\mathcal{P} \rightarrow L} = \frac{1}{\beta} \operatorname{arctanh} \left[\tanh \beta \left(\sum_i^{2(D-1)-1} U_{\mathcal{U}_i \rightarrow \mathcal{U}} + J_{\mathcal{U}} \right) \right. \\
 \left. \tanh \beta \left(\sum_i^{2(D-1)-1} U_{\mathcal{R}_i \rightarrow \mathcal{R}} + J_{\mathcal{R}} \right) \tanh \beta \left(\sum_i^{2(D-1)-1} U_{\mathcal{D}_i \rightarrow \mathcal{D}} + J_{\mathcal{D}} \right) \right], \tag{20}$$

where \mathcal{U}_i (resp. \mathcal{R}_i and \mathcal{D}_i) are the $2(D-1)-1$ plaquettes containing the link U (resp. R and D) excluding plaquette \mathcal{P} .

In the high temperature phase, this Dual approximation with all $u = 0$ should be still a valid approach for any dimensionality D . At low temperatures, however, the EA model in more than two dimensions have a spin glass phase transition and, therefore, we expect the Dual approximation to become poorer, as it can not account for a non trivial order parameter.

By running the Dual algorithm for the 3D EA model we have found a divergence of U -fields around $\beta \simeq 0.39$ for bimodal couplings and around $\beta \simeq 0.41$ for Gaussian couplings. This divergence is due to the fact the U -fields get too much self-reinforced under iteration. This divergence does not come as a surprise given that it happens also when studying the simpler pure ferromagnetic Ising model. However in the ferromagnetic model the temperature at which U -fields diverge is always below the critical temperature and so the Dual algorithm still provides a very good description of the entire paramagnetic phase.

Unfortunately in the 3D EA the divergence of U -fields takes place well above the critical temperature (which is $T_c \simeq 1.12$ for bimodal coupling and at $T_c \simeq 0.95$ for Gaussian couplings, see Ref. 18 for a summary of critical temperatures in 3D spin glasses) and this would make the Dual algorithm of very little use. We have studied the origin of this divergence and we have found a general principle for reducing the divergence of U -fields due to self-reinforcement, thus improving the convergence properties of the Dual algorithm. The idea is the following. When writing the Dual approximation as a constraint satisfaction problem with a non uniform prior, see Eq. (15), the constraints may be redundant. This is the case for the 3D cubic lattice: indeed both the number of links (i.e. variables in the dual problem) and the number of plaquettes (i.e. constraints in the dual problem) are $3N$. So, if constraints were independent, the entropy would be null at $\beta = 0$ and negative for $\beta > 0$ (and this is clearly absurd). The solution to the apparent paradox is that constraints are not independent: actually only $2/3$ of these are independent, and the remaining third is uniquely fixed by the value of the former. In this way the correct entropy is recovered at $\beta = 0$, given that a problem with $3N$ unbiased binary variables subject to $2N$ independent parity-check constraints has entropy $N \log(2)$. The dependence among constraints can be easily appreciated by looking at the 6 plaquette around a cube: if 5 of the 6 constraints are satisfied, then the sixth one is automatically satisfied and redundant.

The general rule for improving the convergence of the Dual algorithm is to remove redundant constraints (this principle is similar to the maxent-normal property of region based free energy approximations [6]). Redundant constraints have no role in determining the fixed point values for the beliefs (since they are redundant), but during the iterations they provide larger fluctuations to messages and may be responsible for the lack of convergence. In practice, on a 3D cubic lattice, we may remove redundant constraints in many different ways: the basic rule states that one constraint (i.e. a plaquette) should be removed for each elementary cube, otherwise if a cube remains with its 6 plaquettes at least one redundant constraint will exist. We have used two different ways of removing one constraint per cube and we have found the same results in the entire paramagnetic phase. So, for simplicity, we are going to present data obtained by removing all constraints corresponding to plaquettes in the xy plane.

The Dual algorithm for the 3D EA model on the cubic lattice with no redundant constraints converges for any temperature above $T \simeq 0.8$ and so we can use it to study the entire paramagnetic phase. The lack of convergence deep in the spin glass phase is to be expected. Just as in the 2D case, the Dual algorithm (when converges) still finds the same solution obtained by a Double Loop algorithm, and again, it finds the solution nearly 100 times faster (see Fig. 10). Double Loop has the apparent advantage of converging at any temperature even at very low ones. However, deep in the spin glass phase, where the underlying paramagnetic approximation is clearly inaccurate, we believe that an algorithm (like the Dual one) that stops converging, is providing an important




FIG. 10. Running times of the Double Loop algorithm [7, 16] (libdai) and of the Dual algorithm on a $8 \times 8 \times 8$ EA model with bimodal interactions ($J_{ij} = \pm 1$). The Dual algorithm is generally several orders of magnitude faster and returns the same solution as the Double Loop algorithm.

warning that something wrong is probably happening. Such a warning would be lacking by using a Double Loop algorithm.

In Figure 11 the correlations predicted by the Dual approximation, and those obtained by a Parallel Tempering Monte Carlo simulation are compared. At high temperatures the correspondence is quite good, but not as good as in 2D. However, it is important to stress that the 3D EA model is much more difficult to simulate than the 2D case: there is no exact method for computing the thermodynamics (at variance to the 2D case) and Monte Carlo methods require huge thermalization times, while the Dual algorithm runs in *linear* time with the system size.

In Figure 12 we show the estimates for the energy obtained from the Monte Carlo and the Dual algorithm (both on a single sample and on the average case). The very strong agreement between the Dual algorithm results on single samples and on the average case is telling us that U -messages arriving at a given point on the lattice are uncorrelated to a very large extent. In other words, the effect of short loops in the lattice is not manifestly present in correlations between messages. On the contrary, the comparison between Dual algorithm results and Monte Carlo results is good only at high temperatures, and it degrades when approaching the critical temperature. This discrepancy can be understood as due to a growing correlation length in the EA model that diverges at the critical temperature: our Dual approximation does not account for correlations beyond the plaquette level and so it becomes inevitably poorer when the correlation length diverges. However, given the extremely fast converging times of the Dual algorithm, it can be viewed as a very effective algorithm for sampling the high temperature paramagnetic phase and as a reasonable approximate algorithm when approaching the critical point.




FIG. 11. Comparison between the correlations $\langle s_i s_j \rangle_{\text{Dual}}$ obtained with the Dual algorithm and the (nearly) exact correlations $\langle s_i s_j \rangle_{\text{PT}}$ obtained with a Parallel Tempering simulation in a 3D EA model of size $8 \times 8 \times 8$ with random bimodal interactions $J_{ij} = \pm 1$. At each temperature the correlation coefficient ρ is reported. For the lowest temperature shown, $\beta = 1.0$, we also report the fraction $w-s$ of pairs of spins such that $\langle s_i s_j \rangle_{\text{Dual}} \langle s_i s_j \rangle_{\text{PT}} < 0$.

VI. CONCLUSIONS

We have introduced a novel Dual algorithm to compute marginals probabilities in the paramagnetic phase of frustrated spin models (e.g. spin glasses) on finite dimensional lattices. Inspired by the fact that in a paramagnetic phase with no external field each variable is unbiased (i.e. local magnetizations are null), the Dual algorithm is derived by adding such paramagnetic constraints in the GBP equations. While BP (i.e. Bethe approximation) and GBP algorithms have serious convergence problems at low temperatures even in the paramagnetic phase, the Dual algorithm converges very fast in a much wider range thanks to these constraints. The Dual algorithm can also be seen as BP on the dual lattice, where the interactions J_{ij} act as external fields on dual variables, thus improving convergence properties of the message passing algorithm.

We have tested the Dual algorithm for the Edwards Anderson spin glass model with bimodal and Gaussian couplings on 2D (square) and 3D (cubic) lattices. The results are very encouraging, showing convergence in the whole paramagnetic phase (and even slightly in the frozen phase for the 3D EA model) and comparing very well with exact correlations measured in Monte Carlo simulations. A comparison with a Double Loop algorithm (which is the state-of-the-art among general purpose inference algorithms) shows that both algorithms found the same result, but our Dual algorithm runs roughly 100 times faster. We also tried to push the Dual approximation to the limit, and we used the correlations inferred by the Dual algorithm to compute ground states configurations in the 2D EA model by a freezing procedure. Again, we showed that the ground states obtained in this way compare very well with exact computations.

The success of our proposal clearly shows that as long as variables are not long range correlated,




FIG. 12. The energy predicted by the Dual approximation in 3D EA model, compared to the average case energy, and the Monte Carlo simulation. We used a $8 \times 8 \times 8$ system with both types of random interactions, bimodal ($J_{ij} = \pm 1$) and Gaussian distributed.

the computation of correlations in a generic spin model can be done in a very fast way by means of message passing algorithms, based on mean-field like approximations. This kind of inference algorithms do not provide in general an exact answer (unless one uses it at very high temperatures or on locally tree-like topologies), and so they can not be seen as substitutes for a Monte Carlo (MC) sampling. However there are many situations where a fast and approximate answer is required more than a slow and exact answer. Let us just make a couple of examples of these situations. On the one side, if one needs to sample from very noisy data, an approximated inference algorithm whose level of approximation is smaller than data uncertainty is as valid as a perfect MC sampler. On the other side, if one needs to use the inferred correlations as input for a second algorithm (as for the freezing algorithm in Section IV B) that will eventually modify/correct these correlations, a fast and reasonably good inference is enough.

The promising results shown in the present work naturally ask for an improvement in several directions. For example, in the paramagnetic phase of a model defined on a 3D lattice, our inference algorithm could be improved by using the $2 \times 2 \times 2$ cube as the elementary region, instead of the plaquette. An even more important improvement would be to extend the applicability range of the algorithm to the low temperature phase: but this requires a rather non trivial modification, since in low temperatures phase the assumption of zero local magnetizations needs to be broken.

ACKNOWLEDGMENTS

F. Ricci-Tersenghi acknowledges financial support by the Italian Research Minister through the FIRB project RBFR086NN1 on “Inference and optimization in complex systems: from the

thermodynamics of spin glasses to message passing algorithms”.

-
- [1] M. Mézard and A. Montanari, *Information, physics, and computation*, Oxford University Press (2009).
 - [2] A. Pelizzola, J. Phys. A **38**, R309 (2005).
 - [3] A.P. Young ed., *Spin glasses and random fields*, World Scientific (1998).
 - [4] Y. Kabashima and D. Saad, Europhys. Lett. **44**, 668 (1998).
 - [5] R. Kikuchi, Phys. Rev. **81**, 988 (1951).
 - [6] J. Yedidia, W. T. Freeman, and Y. Weiss, IEEE Transactions on Information Theory **51**, 2282 (2005).
 - [7] T. Heskes, C. A. Albers, and H. J. Kappen, Proceedings of UAI-2003, 313 (2003).
 - [8] Y. Kabashima, J. Phys. Soc. Jpn. **74**, 2133 (2005).
 - [9] T. Rizzo, A. Lage-Castellanos, R. Mulet, and F. Ricci-Tersenghi, J. Stat. Phys. **139**, 375 (2010).
 - [10] M. Mézard and G. Parisi, Eur. Phys. J. B **20**, 217 (2001); J. Stat. Phys. **111**, 1 (2003). T. Castellani, F. Krzakala, and F. Ricci-Tersenghi, Eur. Phys. J. B **47**, 99 (2005).
 - [11] T. Jorg, J. Lukic, E. Marinari, and O. Martin, Phys. Rev. Lett. **96**, 237205 (2006).
 - [12] M. Mézard and R. Zecchina, Phys. Rev. E **66**, 056126 (2002). A. Montanari, F. Ricci-Tersenghi, and G. Semerjian, J. Stat. Mech. P04004 (2008). F. Ricci-Tersenghi and G. Semerjian, J. Stat. Mech. P09001 (2009).
 - [13] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, Phys. Rev. Lett. **89**, 268701 (2002). L. Zdeborova and F. Krzakala, Phys. Rev. E **76**, 031131 (2007).
 - [14] J. M. Mooij and H. J. Kappen, IEEE Transactions on Information Theory **53**, 4422 (2007).
 - [15] Y. S.-K. Eye and A. L. Yuille, Neural Computation **14**, 2002 (2001).
 - [16] J. M. Mooij, Journal of Machine Learning Research **11**, 2169 (2010). <http://www.libdai.org/>
 - [17] <http://www.informatik.uni-koeln.de/lj-juenger/research/sgs/index.html>
 - [18] H. G. Katzgraber, M. Koerner, and A. P. Young, Phys. Rev. B **73**, 224432 (2006).